

WLCG

Minutes of the 4th Collaboration Board Meeting

(Held at Prague on 22 March 2009)

Present:

CERN IT Head
 CERN PH Deputy Head
 Scientific Secretary
 LCG Project Leader

J. Shiers (for F. Hemmer)
 L. Mapelli
 J. de Groot
 I. Bird

LHC Experiment Spokespersons:

ATLAS Spokesperson
 CMS Spokesperson
 LHCb Spokesperson

D. Barberis (for F. Gianotti)
 M. Kasemann (for J. Virdee)
 Ph. Charpentier (for A. Golutvin)

International Membership

Australia, University of Melbourne
 Canada, TRIUMF, CB Chair
 CERN-IT Grid Support
 China - Tier2, IHEP Beijing
 Czech Rep., FZU AS, Prague
 Finland, NDGF/HIP Tier-2
 France, CC-IN2P3
 France, GRIF, Paris
 Germany, ATLAS Federation, Munich
 Germany, DESY, Hamburg
 Germany, GSI, Darmstadt
 Israel, ICHEP
 Italy, CNAF
 Italy, INFN ATLAS Federation
 Japan, ICEPP, Tokyo
 Netherlands, LHC/Tier1
 Nordic Data Grid Facility (NDGF)
 Norway, UNINETT SIGMA Tier-2
 Poland, Polish Tier-2 Federation
 Romania, Romanian Tier-2 Federation
 Spain, ATLAS Federation
 Spain, LHCb Federation

Spain, PIC
 Sweden, SNIC Tier-2
 Switzerland, CHIP
 Turkey, Turkish Tier-2 Federation (ULAKBIM)
 UK, London Tier-2
 UK, NorthGrid
 UK, RAL
 USA, Caltech CMS T2
 USA, FNAL
 USA, Great Lakes ATLAS T2 Federation

T. Dyce
 M. Vetterli
 J. Shiers
 G. Chen
 M. Lokajicek
 P. Eerola (for D. Riska)
 F. Hernandez (for F. Malek)
 J. Meyer
 G. Dukeck (for S. Bethke)
 V. Guelzow
 P. Malzacher
 L. Levinson
 L. Perini (for M. Mazzucato)
 G. Carlino
 H. Sakamoto
 J. Templon
 L. Fischer
 M. Gronager (for J. Koster)
 M. Witek (for R. Gokieli)
 M. Dulea
 J. Salt
 R. Graciani Diaz
 G. Merino (for M. Delfino
 Reznicek)
 O. Smirnova (for S. Holmgren)
 M.-C. Sawley (for C. Grab)
 E. Akkoyun (for B. Ortakaya)
 D. Colling
 R. Jones
 N. Geddes
 H. Newman
 I. Fisk (for V. White)
 S. Mckee

USA, Purdue CMS T2
USA, Southwest ATLAS T2 Federation

N. Neumeister
K. De

Invited:

Chair, Computing Resources Scrutiny Group
 TRIUMF Tier1 Manager
 CERN-IT, ATLAS

D. Espriu
 R. Tarifout
 D. Van Der Ster

Absent:

ALICE Spokesperson
 Austria, Austrian Tier-2 Federation
 Austria, Austrian Tier-2 Federation
 Belgium, Belgian Tier-2 Federation
 Belgium, Belgian Tier-2 Federation (UA, Antwerpen)
 Belgium, Belgian Tier-2 Federation (UCL, Louvain-la-Neuve)
 Canada, Eastern Tier-2 Federation
 CERN DRSC
 Estonia, NICPB
 France, LAPP, Annecy
 France, LPC, Clermont-Ferrand
 France, SUBATECH, Nantes
 Germany, ATLAS Federation, FR/W
 Germany, ATLAS Federation, FR/W
 Germany, FZK-GridKa
 Hungary, HGCC Federation
 Hungary, HGCC Federation (KFKI-RMKI, Budapest)
 Hungary, HGCC Federation (SzTAKI, Budapest)
 India, TIFR, Mumbai
 India, VECC/SINP, Kolkata
 Italy, INFN ALICE Federation
 Italy, INFN CMS Federation
 Italy, INFN LHCb Federation
 Pakistan, Pakistan Tier-2 Federation
 Republic of Korea, KISTI, Daejeon
 Russian Fed., Russian Data-Intensive GRID (RDIG)
 Slovenia, SIGNET
 Spain, CMS Federation
 Taipei, ASGC
 Turkey, Turkish Tier-2 Federation (TAEK)
 Turkey, Turkish Tier-2 Federation (TAEK)
 Turkey, Turkish Tier-2 Federation (ULAKBIM)
 UK, RAL
 UK, Scotgrid
 UK, SouthGrid
 Ukraine, Ukrainian Tier-2 Federation
 USA, BNL
 USA, Florida CMS T2
 USA, Midwest ATLAS T2 Federation
 USA, MIT CMS T2
 USA, Nebraska CMS T2
 USA, Northeast ATLAS T2 Federation
 USA, SLAC ATLAS T2
 USA, U. Wisconsin CMS T2
 USA, UC San Diego CMS T2

A. Aimar
 Y. Cheng
 S. De Weirdt
 J. Schukraft
 D. Kuhn
 C. Wulz
 G. Bruno
 O. Devroede
 P. Vanlaer
 P. Savard
 S. Bertolucci
 M. Kadastik
 S. Jezequel
 D. Pallin
 L. Aphecetche
 T. Hareberg
 J. Sundermann
 K. Mickel
 G. Vesztergombi
 D. Horvath
 C. Hajdu
 A. Gurtu
 Y. Viyogi
 M. Masera
 M. Paganoni
 U. Marconi
 H. Hoorani
 S. Hwang
 V. Ilyin
 B. Kersevan
 F. Matorras
 S. Lin
 L. Baskus
 • Cakir
 B. Akcan
 J. Gordon (was present)
 N. Glover
 P. Watkins
 G. Zinovjev
 M. Ernst
 P. Avery
 R. Gardner (was present)
 C. Paus
 K. Bloom
 J. Shank
 W. Yang
 S. Dasu
 F. Wuerthwein

Introduction

M. Vetterli welcomes the participants to the meeting.

He thanks M. Lokajicek for hosting the meeting. He thanks D. Jacobs for his support for the work of the Collaboration Board as Scientific Secretary. J. de Groot is taking over as Scientific Secretary as of this meeting.

Agenda

The Collaboration Board agrees on the following Agenda:

- Minutes from the last meeting
 - Business arising / Collaboration matters
 - Status of WLCG
 - Report from the Scrutiny Group
 - User Support Information Exchange
 - AOB
 - Next meeting
-

Minutes of the last Meeting

The minutes of the last meeting are approved.

Matters Arising (M. Vetterli)

The question of pledged versus used capacities will be addressed by I. Bird in his presentation. An automated system to gather the data is being developed and nearing completion. A 'manual' survey of Tier2 sites is under way. M. Vetterli urges the Tier2s to respond. I. Bird will address the transition from EGEE to EGI in his presentation.

M. Vetterli summarizes the present situation concerning Tier2 communications.

ALICE designated a Core Offline person to have privileged contact with a given T2 site manager. Tier-2 federations provide a single contact person.

ATLAS uses its cloud structure for communications. Each cloud has a Tier-2 coordinator.

CMS has two Tier-2 coordinators who feed T2 issues back to the operations group.

LHCb uses GGUS tickets to communicate with Tier2s. This works well because LHCb use the Tier2s only for Monte-Carlo.

In addition to these VO specific structures there are regional organizations, notably in France, Canada, the UK and USA.

M. Vetterli further mentions that M. Jouvin and A. Gurtu were appointed Tier2 representatives in the Overview Board. A Tier2 mailing list exists. The Grid Deployment Board in October was dedicated to Tier2 issues. Given the success he suggests that this could be repeated.

WLCG Status Report (I. Bird)

I. Bird mentions that the CCRC '08, performed since the last meeting of the Collaboration Board, was successful, although not everything was fully tested.

Two LHCC mini reviews have taken place. The LHCC made a number of recommendations:

- perform a CCRC '09 - this is being looked into;
- make sure experiments are not limited by resources when data-taking starts;
- an official 2009/2010 schedule is required urgently so all experiments can plan their resources requirements in a consistent manner.

It was noted that the experiments still suffer from SRM (MSS) problems. In the applications area "Very good progress on all fronts with very mature organization, well managed, giving

results". Since then, the experiments have been ready for data taking and took cosmics and some single beam data. Services have improved since.

All Tier1 and Tier2 centres have now signed the WLCG MoU. A new MoU has been signed with the Republic of Korea as a Tier2 centre supporting CMS. Brazil is still planning to sign as a Tier2 centre supporting all four experiments. In Summary, 33 countries have now signed (~50 signatures) for 11 Tier1s, 61 Tier2 federations (120 Tier2 sites).

Missing resources are due to a number of factors: procurement problems, faulty hardware delivered, as well as cooling and power problems. In this context it has been agreed with the C-RRB to advance the annual planning cycle to take into account uncompressible delays in procurement.

I. Bird shows statistics on throughput, usage patterns and reliabilities. Reliabilities for Tier1 and Tier2 centres have been improving. Unavailability of Tier1 services is caused not by a single but by multiple factors. A monthly availability/reliability report is compiled. The target is to get all Tier2 sites above 95%. A VO-specific report on Tier1 reliabilities is in the process of being validated.

The pledge balance as presented to the RRB meeting in November is shown. Nothing has changed since. In particular, the shortfall of resources for ALICE remains. The RRB has agreed to a reduction of the planning 'window' from five to three years.

The presently agreed LHC schedule implies continuous running for up to 44 weeks from ~November 2009 with only a short break for the year-end holidays. The LHC energy will be limited to 2*5 TeV. A heavy ion run is planned at the end of this period. The experiments are re-assessing their requirements in light of an effective amount of data in 2009 – 2010 corresponding to about $6 \times 10^{**6}$ seconds. The deadline is March 31 so that consolidated data can be provided to the RRB meetings in April.

A new performance benchmark (HEP-SPEC06) has been agreed. Sites must benchmark their existing capacity and vendors will be asked to run this benchmark suite. A procedure for the automatic gathering of data on of installed capacity has been agreed by EGEE and OSG and will be implemented.

A large-scale combined test in 2009 did not fit with the experiments' original schedules. But the Tier1s are concerned that there has not been a combined test with all experiments testing tape recall/reprocessing at nominal rates. Large-scale tests of analysis have not been performed. Furthermore, the LHCC mini-review recommended a CCRC'09 in some form. There was agreement on this in the workshop preceding the CB meeting. The tests will take place in May and June.

The 2009 capacity should be fully commissioned in October and the 2010 capacity in April.

Visits to Tier1 sites are planned for 2009 to understand better a number of service issues.

The EGEE project comes to an end in April 2010. A final draft of the EGI blueprint has been produced. At the EGI Policy Board meeting in Catania, Amsterdam was chosen as the location for EGI.org. The WLCG Overview Board has sent a letter of support for the EGI/NGI initiative to the chair of the EGI Policy Board.

Discussion:

The transfer from EGEE to EGI will dominate the planning for 2010. There are proposals to be made for the upcoming funding calls. J. Templon says that the Amsterdam group will find funds to set up a team by summer. Legal structures have to be put into place and a MoU will have to be signed. Help is welcome. M. Jouvin adds that a temporary team is needed to produce funding proposals. He sees good progress in this area.

I. Bird notes that it is crucial for WLCG to ensure continuity at the present stage of the LHC program. In his view it will be necessary to take control of the middleware. LCG may have to distribute the software. The alternative would be a gLite consortium. He adds that WLCG needs to maintain a working system for the LHC experiments.

J. Templon notes that owning the software may hamper the funding from national funding sources that fund more than just HEP.

M. Vetterli concludes by saying that the momentum appears to be in the direction of owning the middleware. This clearly needs a wider discussion.

I. Bird concludes by saying that CCRC'08 was successful and that the service has continued to be used at significant levels. The ramp-up of resources for 2009/10 continues.

Report from C-RSG Chair

D. Espriu, Chair of the Computing Resources Scrutiny Group, thanks M. Vetterli for the opportunity to explain the work of the C-RSG.

The mission and composition of the C-RSG are defined in the MoU. The task of the C-RSG is to scrutinize the requests for and usage of computing resources by the experiments. The objective is to check if the computing model is implemented correctly, not to validate the computing model itself. The latter is the responsibility of the LHCC. He shows the composition of the group and emphasizes the independence of its members.

The C-RSG reported to the RRB meeting in November 2008. Both the 2008 and 2009 resources requests had been scrutinized. D. Espriu thanks the experiments for their collaborative spirit and openness.

The C-RSG based their work on the experiments' requests submitted in Sept. 2007, modified in the case of ATLAS. The 19 September incident forced a change of the running scenario initially considered. The updated scenario, the basis for the report to the C-RRB in November, assumed 2009 to be a nearly standard running year with $0.9 \cdot 10^{*7}$ s. of pp and 10^{*6} s. AA running.

The C-RSG wishes to bring a number of issues to the attention of the LHCC:

- Most experiments propose using increased trigger rates as compared to the ones stated in the TDR, reviewed by the LHCC.
- ALICE wants to increase substantially their amount of pp data.
- ALICE should undertake a full assessment of the impact on physics if requested resources do not materialize.
- CMS have made an effort to reduce the event sizes. This example should be followed by all experiments.
- The C-RSG takes note of potential modifications of the computing models due to the use of different data formats serving the same purposes, not always well justified.
- The implementation of the ATLAS computing model differs slightly from the one envisaged in the TDR. This implies justified increased demands on CERN resources.
- It seems questionable to support substantial requests based on cosmic runs, but the C-RSG does not have sufficient insight to make a definite scientific judgement.

The C-RSG makes the following general comments and recommendations:

- It seems prudent to scrutinise the experiments' use of resources after a few months of data taking in 2009. Given the resource acquisition cycle, the Tier1 and Tier2 centres should be informed of the 2010 resource acquisition plans as soon as possible. The C-RSG will provide a scrutiny at the earliest feasible date.
- It would be very helpful to the funding agencies and institutes to have a scrutiny ready by summer, thus giving more time to the Tier1 and Tier2 for the procurement process for the following year.
- In view of the unprecedented scale of the LHC computing effort the C-RSG recommends that the collaborations undertake a risk analysis and take into account their results in future requests in order to cope with the most likely failures or shortfalls.
- The information provided about the AA program was sometimes rather sketchy. The C-RSG would be thankful for more details in the future.
- Changing running conditions (75ns bunch crossings) will have to be accommodated within the existing envelope by decreasing the event rate or similar measures.

- Event sizes should be reduced as much as possible as detectors become better understood.
- A very strict policy of removing all 'dark' or 'orphaned' data should be enforced.
- Experiments should keep their computing models and needs under constant revision. Some requests remained unchanged even though no longer realistic.
- The experiments should make maximal use of the resources on the grid for analysis, avoiding as much as possible the use of CERN facilities.
- A clear separation is advocated between the use of CERN resources for calibration, first pass reconstruction and central analysis ('express stream' or similar), and those used to perform physics analysis by CERN based physicists.
- The C-RSG wishes to state that the recommendations contained in the scrutiny report are to the best of its knowledge rigorous. They correspond to the real needs of the experiments for a given LHC live time in the present stage of the commissioning and of their computing model implementation. Shortfalls would seriously jeopardize the success of the experiments. The C-RSG therefore recommends that the funding agencies ensure the effective and timely delivery of the pledged resources.

Conclusions:

- In the process of scrutinizing the 2008 and 2009 requests of the four LHC experiments the C-RSG has critically examined all possible aspects of the different computing models and their implementation.
- While some points of discrepancy and a few potentially troublesome issues exist, the overall demand of resources for 2009 largely remains within the envisaged envelope.
- A very limited degree of redistribution of resources may be advisable in 2009, however care has to be taken not to harm experiments with a more consolidated CM in favour of those whose CM is less defined or consolidated at this stage.
- To remain in the future within the envelope will require some updates and revisions of the computing models, perhaps of some substance in some cases. The scrutiny after the first round of real data will be of great relevance.
- The CRSG believes that the different computing models have largely proven their validity and has no doubt that they will survive their first contact with real data in 2009.

User Support (D. van der Ster, R. Tafirout)

D. van der Ster (CERN-IT Grid Support & ATLAS) provides some details on Grid User Support in ATLAS.

The ATLAS model is based on regular tutorials and a help forum. ATLAS formed a Distributed Analysis Support Team (DAST). The goal was to provide support for the Pathena and Ganga analysis tools thereby relieving developers of support tasks. DAST consists of a team of volunteer expert shifters. The service started in October last year and currently provides one expert on shift each for the European and American time zones with Asia-Pacific missing at present. DAST is not a traditional help desk. The support uses an eGroups forum to enable user to user support.

D. van der Ster recounts some of the experience gained so far. User queries show an expected pattern but indicate that more than just DA support is required. Many users are not aware of existing tools. Conclusion is that end-user tools need to be fully integrated with monitoring and should enforce policies.

For the future a number of weaknesses in documentation and tools need to be addressed to provide more user-oriented monitoring and 'status awareness' in the end-user tools. Users are asking questions already covered in existing documentation indicating the need for improvement the area of tutorials, twiki documents etc.

ATLAS estimates that the manpower for DAST shifters needs to be doubled to provide 2 shifters each for Europe, the Americas and Asia.

R. Tafirout (TRIUMF) presents some aspects of user support from the Tier1 viewpoint. Two staff were hired at TRIUMF in 2007 to provide user support. The support staff interact both with the Tier1 and Tier2 operations teams and primarily target the Canadian cloud users. The Canadian cloud supports ATLAS only.

Documentation and howto's are available on the CA twiki. New releases of ATLAS software or patches are validated before installation. Tier1 and Tier2 outputs are compared.

The experience so far indicates that Tier1 and Tier2s have been fully exercised. Problems with the PANDA and GANGA systems are fed back to the systems developers. Support is also provided for Tier3 analysis testing. Typical user issues concern data management and replication issues, ATLAS specific and storage issues.

Discussion:

Ph. Charpentier comments that this concerns ATLAS rather than general support. M. Vetterli agrees but notes that other VOs would be supported as well if Canada were involved by hiring more people.

M. Vetterli raises the problem of the size of FAQs and Twikis. These are becoming very large and difficult to handle. Diagnostics are insufficient.

D. van der Ster confirms that better diagnostics are necessary; it is difficult to reproduce errors. On a question he replies that the tools developed are not Canada-specific and could be generalized in the ATLAS context.

Closing Remarks (M. Vetterli)

The transfer from EGEE to EGI needs a much larger discussion.

With reduced manpower, in particular also due to the transfer from EGEE to EGI, the WLCG Collaboration needs to become more pro-active. The dCache workshop was user-driven and could serve as a model for a similar effort in the area of data storage. This needs further discussion.

The subject of Service Credits for Common Computing needs to be addressed. Some experiments currently practice it, but not others. There should be a consistent approach. Care has to be taken not to double-count those people supporting more than one activity. WLCG could verify the claims. M. Vetterli intends to discuss this with the experiments.

A short discussion on Service Credits shows that there are indeed different approaches and that opinions differ.

Next Meeting

A second meeting of the WLCG Collaboration Board in 2009 is planned for the fall, probably in conjunction with a workshop in order to minimize overhead.
